**El propósito de este documento es levantar una red VXLAN clásica (LEAF/SPINE) usando multicast como replicación.**

# Topología

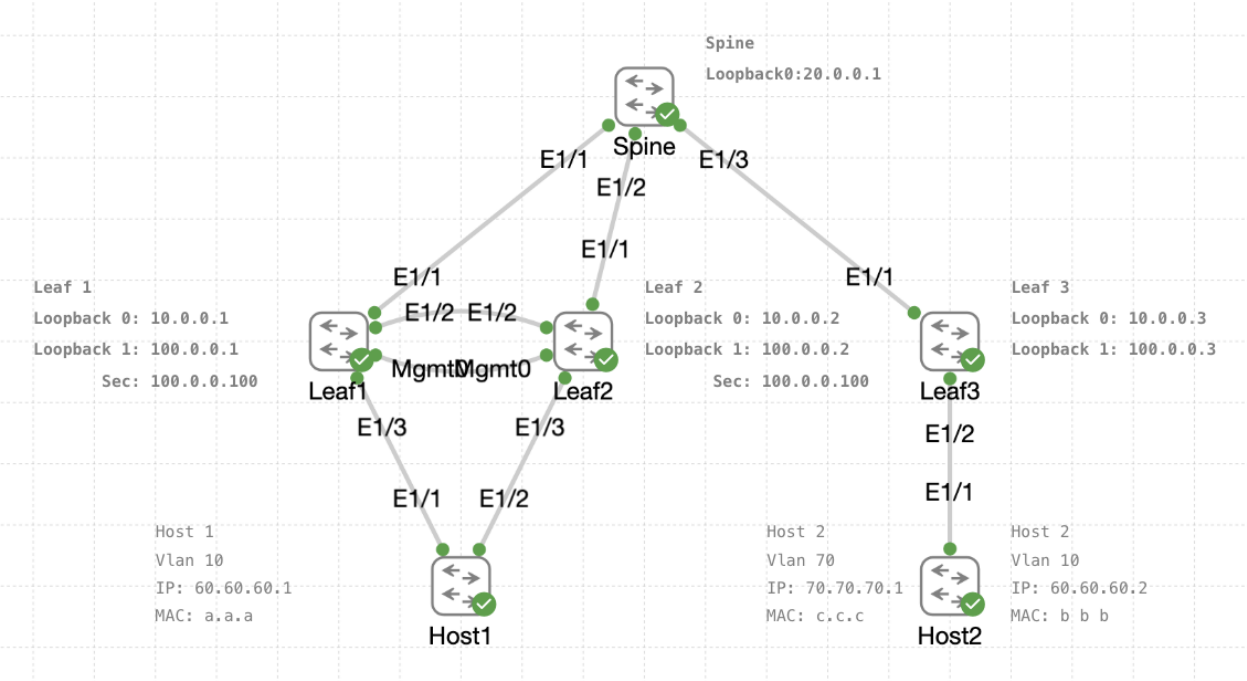


# VPC

### Leaf 1

```
interface mgmt0
  ip address 1.1.1.1/24
  no shutdown

feature vpc
feature lacp

vpc domain 1
  peer-keepalive destination 1.1.1.2 source 1.1.1.1


interface Ethernet1/2
  channel-group 1 mode active
```

```
interface port-channel1
switchport
  switchport mode trunk
  vpc peer-link
  no shutdown

interface Ethernet1/3
  channel-group 2 mode active


interface port-channel2
  switchport
  switchport mode trunk
  vpc 2
  no shutdown
```

**Leaf 2**

```
interface mgmt0
  ip address 1.1.1.2/24
  no shutdown

feature vpc
feature lacp

vpc domain 1
  peer-keepalive destination 1.1.1.1 source 1.1.1.2


interface Ethernet1/2
  channel-group 1 mode active

interface port-channel1
  switchport
  switchport mode trunk
  vpc peer-link
  no shutdown

interface Ethernet1/3
  channel-group 2 mode active


interface port-channel2
  switchport
  switchport mode trunk
  vpc 2
  no shutdown
```

### Verificación:

```
sh vpc
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 1
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : primary
Number of vPCs configured         : 1
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Disabled
Delay-restore status              : Timer is off.(timeout = 30s)
Delay-restore SVI status          : Timer is off.(timeout = 10s)
Operational Layer3 Peer-router    : Disabled
Virtual-peerlink mode             : Disabled

vPC Peer-link status
---------------------------------------------------------------------
id    Port   Status Active vlans
--    ----   ------ --------------------------------------------------
1     Po1    up     1


vPC status
----------------------------------------------------------------------------
Id    Port          Status Consistency Reason                Active vlans
--    ------------  ------ ----------- ------                ---------------
2     Po2           up     success     success               1
```

# Underlay

### Leaf 1

```
interface loopback0
  description underlay_loopback
  ip address 10.0.0.1/32

interface loopback1
  description overlay_loopback
  ip address 100.0.0.1/32
  ip address 100.0.0.100/32 secondary

interface Ethernet1/1
  description uplink_to_spine
  no switchport
  medium p2p
  ip unnumbered loopback0
  no shutdown
```

**Leaf 2**

```
interface loopback0
  description underlay_loopback
  ip address 10.0.0.2/32

interface loopback1
  description overlay_loopback
  ip address 100.0.0.2/32
  ip address 100.0.0.100/32 secondary

interface Ethernet1/1
  description uplink_to_spine
  no switchport
  medium p2p
  ip unnumbered loopback0
  no shutdown
```

**Leaf 3**

```
interface loopback0
  description underlay_loopback
  ip address 10.0.0.4/32

interface loopback1
  description overlay_loopback
  ip address 100.0.0.4/32

interface Ethernet1/1
  description uplink_to_spine
  no switchport
  medium p2p
  ip unnumbered loopback0
  no shutdown
```

**Spine**

```
interface loopback0
  description underlay_loopback
  ip address 20.0.0.1/32

interface Ethernet1/1-3
description downlink_to_leaf
  no switchport
  medium p2p
  ip unnumbered loopback0
  no shutdown
```

### Verificación:

Todas las interfaces tienen que estar en **protocol-up/link-up/admin-up**

**Leaf1# sh ip int brief**

```
IP Interface Status for VRF "default"(1)
Interface            IP Address      Interface Status
Lo0                  10.0.0.1        protocol-up/link-up/admin-up
Lo1                  100.0.0.1       protocol-up/link-up/admin-up
Eth1/1               unnumbered      protocol-up/link-up/admin-up
                     (loopback0)
```

**Leaf2# sh ip int brief**

```
IP Interface Status for VRF "default"(1)
Interface            IP Address      Interface Status
Lo0                  10.0.0.2        protocol-up/link-up/admin-up
Lo1                  100.0.0.2       protocol-up/link-up/admin-up
Eth1/1               unnumbered      protocol-up/link-up/admin-up
                     (loopback0)
```

**Leaf3# sh ip int brief**

```
IP Interface Status for VRF "default"(1)
Interface            IP Address      Interface Status
Lo0                  10.0.0.3        protocol-up/link-up/admin-up
Lo1                  100.0.0.3       protocol-up/link-up/admin-up
Eth1/1               unnumbered      protocol-up/link-up/admin-up
                     (loopback0)
```

**Spine# sh ip int brief**

```
IP Interface Status for VRF "default"(1)
Interface            IP Address      Interface Status
Lo0                  20.0.0.1        protocol-up/link-up/admin-up
Eth1/1               unnumbered      protocol-up/link-up/admin-up
                     (loopback0)
Eth1/2               unnumbered      protocol-up/link-up/admin-up
                     (loopback0)
Eth1/3               unnumbered      protocol-up/link-up/admin-up
                     (loopback0)
```

# OSPF

### Leaf 1-3

```
feature ospf

router ospf 1

interface loopback0
  ip router ospf 1 area 0.0.0.0

interface loopback1
  ip router ospf 1 area 0.0.0.0
```

```
interface Ethernet1/1
  ip router ospf 1 area 0.0.0.0
```

**Spine**

```
feature ospf
router ospf 1

interface loopback0
  ip router ospf 1 area 0.0.0.0

interface Ethernet1/1
  ip router ospf 1 area 0.0.0.0

interface Ethernet1/2
  ip router ospf 1 area 0.0.0.0

interface Ethernet1/3
  ip router ospf 1 area 0.0.0.0
```

**Verificación:**

Si nos movemos al SPINE, se tienen que observar a los 3 leafs como OSPF neighbors

```
Spine# sh ip ospf neighbors
 OSPF Process ID 1 VRF default
 Total number of neighbors: 3
 Neighbor ID     Pri State            Up Time  Address         Interface
 10.0.0.1          1 FULL/ -          2d18h    10.0.0.1        Eth1/1
 10.0.0.2          1 FULL/ -          2d18h    10.0.0.2        Eth1/2
 10.0.0.3          1 FULL/ -          2d18h    10.0.0.3        Eth1/3
```

# PIM

**Leaf 1-3**

```
feature pim
ip pim rp-address 20.0.0.1 group-list 224.0.0.0/4

interface loopback0
  ip pim sparse-mode

interface loopback1
  ip pim sparse-mode

interface Ethernet1/1
  ip pim sparse-mode
```

### Spine

```
feature pim

ip pim rp-address 20.0.0.1 group-list 224.0.0.0/4

interface loopback0
  ip pim sparse-mode

interface Ethernet1/1
  ip pim sparse-mode

interface Ethernet1/2
  ip pim sparse-mode

interface Ethernet1/3
  ip pim sparse-mode
```

### Verificación:

Si nos movemos al SPINE, se tienen que observar a los 3 leafs como PIM neighbors

```
Spine# sh ip pim neighbor
PIM Neighbor Status for VRF "default"
Neighbor        Interface            Uptime    Expires   DR       Bidir-  BFD    ECMP Redirect
                                                         Priority Capable State     Capable
10.0.0.1        Ethernet1/1          2d18h     00:01:36  1        yes     n/a     no
10.0.0.2        Ethernet1/2          2d18h     00:01:31  1        yes     n/a     no
10.0.0.3        Ethernet1/3          2d18h     00:01:41  1        yes     n/a     no
```

# BGP

### Leaf 1-3

```
feature bgp
feature nv overlay
nv overlay evpn

router bgp 1
  address-family l2vpn evpn
  neighbor 20.0.0.1
    remote-as 1
    update-source loopback0
    address-family l2vpn evpn
      send-community both
```

### Spine

```
feature bgp
feature nv overlay
nv overlay evpn
```

```
router bgp 1
  address-family l2vpn evpn

  neighbor 10.0.0.1
    remote-as 1
    update-source loopback0
    address-family l2vpn evpn
      route-reflector-client
      send-community both

  neighbor 10.0.0.2
    remote-as 1
    update-source loopback0
    address-family l2vpn evpn
    route-reflector-client
    send-community both

  neighbor 10.0.0.4
    remote-as 1
    update-source loopback0
    address-family l2vpn evpn
      route-reflector-client
      send-community both
```

**Verificación:**

Si nos movemos al SPINE, se tienen que observar a los 3 leafs como BGP L2VPN neighbors

```
Spine# sh bgp l2vpn  evpn summary
BGP summary information for VRF default, address family L2VPN EVPN
BGP router identifier 20.0.0.1, local AS number 1
BGP table version is 71, L2VPN EVPN config peers 3, capable peers 3
10 network entries and 10 paths using 2440 bytes of memory
BGP attribute entries [9/1548], BGP AS path entries [0/0]
BGP community entries [0/0], BGP clusterlist entries [0/0]

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down
State/PfxRcd
10.0.0.1        4     1    4010    4004       71    0    0 02:55:39 0
10.0.0.2        4     1    4009    4006       71    0    0 02:55:30 0
10.0.0.3        4     1    4004    4008       71    0    0 02:55:20 0
```

# VXLAN Switching

### Leaf 1-3

```
feature vn-segment-vlan-based


vlan 10
      vn-segment 100010


interface nve1
  host-reachability protocol bgp
  source-interface loopback1
  member vni 100010
    mcast-group 239.0.0.10
  no shutdown

evpn
  vni 100010 l2
    rd auto
    route-target both auto
```

### Verificación:

En todos los leafs, revisar en los siguientes comandos que aparezcan las secciones remarcadas

```
sh nve vni
Codes: CP - Control Plane        DP - Data Plane
       UC - Unconfigured         SA - Suppress ARP
       SU - Suppress Unknown Unicast
       Xconn - Crossconnect
       MS-IR - Multisite Ingress Replication

Interface VNI      Multicast-group   State Mode Type [BD/VRF]      Flags
--------- -------- ----------------- ----- ---- ------------------ -----
nve1      100010   239.0.0.10        Up    CP   L2 [10]


sh nve interface nve 1 detail
Interface: nve1, State: Up, encapsulation: VXLAN
 VPC Capability: VPC-VIP-Only [notified]
 Local Router MAC: 5213.093b.1b08
 Host Learning Mode: Control-Plane
 Source-Interface: loopback1 (primary: 100.0.0.1, secondary: 100.0.0.100)
 Source Interface State: Up
 Virtual RMAC Advertisement: No
 NVE Flags:
 Interface Handle: 0x49000001
```

```
 Source Interface hold-down-time: 180
 Source Interface hold-up-time: 30
 Remaining hold-down time: 0 seconds
 Virtual Router MAC: 0200.6400.0064
 Interface state: nve-intf-add-complete
```

**Host 1**

```
Feature lacp
Feature interface-vlan

interface Ethernet1/1-2
  channel-group 2 mode active


interface port-channel2
  switchport
  switchport mode trunk

vlan 10

interface Vlan10
  no shutdown
  mac-address 000a.000a.000a
  ip address 60.60.60.1/24
```

**Host 2**

```
Feature interface vlan

Interface e1/1
  Switchport
  Switchport mode trunk

vlan 10

interface Vlan10
  no shutdown
  mac-address 000b.000b.000b
  ip address 60.60.60.2/24
```

**Leaf3**

```
interface Ethernet1/3
  switchport mode trunk
```

**Verificación de conectividad:**

```
Host1# ping 60.60.60.2
PING 60.60.60.2 (60.60.60.2): 56 data bytes
64 bytes from 60.60.60.2: icmp_seq=0 ttl=254 time=8.794 ms
64 bytes from 60.60.60.2: icmp_seq=1 ttl=254 time=8.22 ms
64 bytes from 60.60.60.2: icmp_seq=2 ttl=254 time=6.158 ms
64 bytes from 60.60.60.2: icmp_seq=3 ttl=254 time=6.715 ms
64 bytes from 60.60.60.2: icmp_seq=4 ttl=254 time=6.849 ms

--- 60.60.60.2 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 6.158/7.347/8.794 ms
```

# VXLAN Routing

**Leaf 1-3**

```
fabric forwarding anycast-gateway-mac e.e.e
feature interface-vlan

vlan 50
  vn-segment 100050

vrf context tenant-1
  vni 100050
  rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn

interface Vlan50
  no shutdown
  vrf member tenant-1
  no ip redirects
  ip forward

interface nve 1
member vni 100050 associate-vrf

route-map direct-bgp

router bgp 1
  vrf tenant-1
    address-family ipv4 unicast
      redistribute direct route-map direct-bgp
```

**Verificación:**

En todos los leafs verificar las secciones remarcadas.

```
sh nve vni
Codes: CP - Control Plane        DP - Data Plane
       UC - Unconfigured         SA - Suppress ARP
       SU - Suppress Unknown Unicast
       Xconn - Crossconnect
       MS-IR - Multisite Ingress Replication

Interface VNI      Multicast-group   State Mode Type [BD/VRF]
Flags
--------- -------- ----------------- ----- ---- ------------------ ---
--
nve1      100010   239.0.0.10        Up    CP   L2 [10]
nve1      100050   n/a               Up    CP   L3 [tenant-1]

sh ip int brief vrf tenant-1

IP Interface Status for VRF "tenant-1"(3)
Interface            IP Address      Interface Status
Vlan50               forward-enabled protocol-up/link-up/admin-up
```

**Leaf 1-2**

```
interface Vlan10
  vrf member tenant-1
  ip address 60.60.60.254/24
  fabric forwarding mode anycast-gateway
  no shutdown
```

**Verificación:**

Verificar las secciones remarcadas.

```
Leaf1# sh ip int brief vrf tenant-1

IP Interface Status for VRF "tenant-1"(3)
Interface            IP Address      Interface Status
Vlan10               60.60.60.254    protocol-up/link-up/admin-up
Vlan50               forward-enabled protocol-up/link-up/admin-up
```

**Leaf 3**

```
vlan 70
  vn-segment 100070

interface Vlan70
  no shutdown
  vrf member tenant-1
  ip address 70.70.70.254/24
  fabric forwarding mode anycast-gateway



interface nve1
  member vni 100070
    mcast-group 239.0.0.70

evpn

  vni 100070 l2
    rd auto
    route-target both auto
```

**Verificación:**

Verificar las secciones remarcadas

```
Leaf3# sh nve vni
Interface VNI      Multicast-group   State Mode Type [BD/VRF]
Flags
--------- -------- ----------------- ----- ---- ----------------
-- -----
nve1      100010   239.0.0.10        Up    CP   L2 [10]
nve1      100050   n/a               Up    CP   L3 [tenant-1]
nve1      100070   239.0.0.70        Up    CP   L2 [70]
```

```
Leaf3# sh ip int br vrf tenant-1

IP Interface Status for VRF "tenant-1"(3)
Interface            IP Address      Interface Status
Vlan50               forward-enabled protocol-up/link-up/admin-up
Vlan70               70.70.70.254    protocol-up/link-up/admin-up
```

**Host 1**

```
ip route  0.0.0.0/0 60.60.60.254
```

**Host 2**

```
Vlan 70

interface Vlan70
  no shutdown
  mac-address 000c.000c.000c
  ip address 70.70.70.1/24

ip route 0.0.0.0/0 70.70.70.254
```

**Verificación:**

```
Host1# ping 70.70.70.1
PING 70.70.70.1 (70.70.70.1): 56 data bytes
64 bytes from 70.70.70.1: icmp_seq=0 ttl=254 time=10.868 ms
64 bytes from 70.70.70.1: icmp_seq=1 ttl=254 time=9.772 ms
64 bytes from 70.70.70.1: icmp_seq=2 ttl=254 time=9.91 ms
64 bytes from 70.70.70.1: icmp_seq=3 ttl=254 time=9.274 ms
64 bytes from 70.70.70.1: icmp_seq=4 ttl=254 time=9.441 ms

--- 70.70.70.1 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 9.274/9.853/10.868 ms
```